

**IEEE Xplore®**  
RELEASE 2.1

Welcome United States Patent and Trademark Office

Home | Login | Logout | Access Information | Alerts

Search Session History      BROWSE      SEARCH      IEEE XPLORE GUIDE

Fri, 2 Dec 2005, 10:50:56 AM EST

Edit an existing query or compose a new query in the Search Query Display.

Select a search number (#) to:

- Add a query to the Search Query Display
- Combine search queries using AND, OR, or NOT
- Delete a search
- Run a search

Search Query Display

Recent Search Queries

#1	((address alignment<in>metadata) <and> (page boundary<in>metadata)) <and> (cache<in>metadata)
#2	((address alignment)<in>metadata)
#3	((address alignment)<in>metadata)
#4	(cache<IN>metadata)
#5	(network adapter<IN>metadata)
#6	(address boundary<IN>metadata)
#7	(end of frame<IN>metadata)

Help Contact Us Privacy  
© Copyright 2005 IEEE

Indexed by  
**Inspec**

 **PORTAL**  
USPTO

Subscribe (Full Service) Register (Limited Service, Free) Login  
 Search:  The ACM Digital Library  The Guide  
 +address +alignment, +page +boundary, +cross +boundary,  SEARCH

The ACM Digital Library  [Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used For 42 167,€  
address alignment page boundary cross boundary cache data transfer network adapter end of frame

Sort results by relevance  Save results to a Binder  Try an Advanced Search  
 Display results expanded form  Search Tips  Try this search in The ACM Guide  
 Open results in a new window

Results 1 - 20 of 42

Result page: 1 2 3 [next](#)Relevance scale **1 Experiences with a high-speed network adaptor: a software perspective**

 Peter Druschel, Larry L. Peterson, Bruce S. Davie

October 1994 **ACM SIGCOMM Computer Communication Review , Proceedings of the conference on Communications architectures, protocols and applications SIGCOMM '94**, Volume 24 Issue 4

Publisher: ACM Press

Full text available:  pdf(1.37 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper describes our experiences, from a software perspective, with the OSIRIS network adaptor. It first identifies the problems we encountered while programming OSIRIS and optimizing network performance, and outlines how we either addressed them in the software, or had to modify the hardware. It then describes the opportunities provided by OSIRIS that we were able to exploit in the host operating system (OS); opportunities that suggested techniques for making the OS more effective in d ...

**2 Fast detection of communication patterns in distributed executions**

Thomas Kunz, Michiel F. H. Seuren

November 1997 **Proceedings of the 1997 conference of the Centre for Advanced Studies on Collaborative research**

Publisher: IBM Press

Full text available:  pdf(4.21 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Understanding distributed applications is a tedious and difficult task. Visualizations based on process-time diagrams are often used to obtain a better understanding of the execution of the application. The visualization tool we use is Poet, an event tracer developed at the University of Waterloo. However, these diagrams are often very complex and do not provide the user with the desired overview of the application. In our experience, such tools display repeated occurrences of non-trivial commun ...

**3 The measured performance of personal computer operating systems**

 J. B. Chen, Y. Endo, K. Chan, D. Mazières, A. Dias, M. Seltzer, M. D. Smith

December 1995 **ACM SIGOPS Operating Systems Review , Proceedings of the fifteenth ACM symposium on Operating systems principles SOSP '95**, Volume 29 Issue 5

Publisher: ACM Press

Full text available:  pdf(1.98 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

4 Performance evaluation and cache analysis of an ILP protocol implementation

Torsten Braun, Christophe Diot

June 1996 **IEEE/ACM Transactions on Networking (TON)**, Volume 4 Issue 3

**Publisher:** IEEE Press

Full text available:  [pdf\(1.57 MB\)](#)

Additional Information: [full citation](#), [references](#), [index terms](#)

5 Workshop on compositional software architectures: workshop report



May 1998 **ACM SIGSOFT Software Engineering Notes**, Volume 23 Issue 3

**Publisher:** ACM Press

Full text available:  [pdf\(2.91 MB\)](#)

Additional Information: [full citation](#), [index terms](#)

6 Securing ATM networks



Shaw-Cheng Chuang

January 1996 **Proceedings of the 3rd ACM conference on Computer and communications security**

**Publisher:** ACM Press

Full text available:  [pdf\(1.53 MB\)](#)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

7 VM-based shared memory on low-latency, remote-memory-access networks



Leonidas Kontothanassis, Galen Hunt, Robert Stets, Nikolaos Hardavellas, Michał Cierniak, Srinivasan Parthasarathy, Wagner Meira, Sandhya Dwarkadas, Michael Scott

May 1997 **ACM SIGARCH Computer Architecture News , Proceedings of the 24th annual international symposium on Computer architecture ISCA '97**, Volume 25 Issue 2

**Publisher:** ACM Press

Full text available:  [pdf\(1.96 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Recent technological advances have produced network interfaces that provide users with very low-latency access to the memory of remote machines. We examine the impact of such networks on the implementation and performance of software DSM. Specifically, we compare two DSM systems---Cashmere and TreadMarks---on a 32-processor DEC Alpha cluster connected by a Memory Channel network. Both Cashmere and TreadMarks use virtual memory to maintain coherence on pages, and both use lazy, multi-writer releases ...

8 The measured performance of personal computer operating systems



J. Bradley Chen, Yasuhiro Endo, Kee Chan, David Mazières, Antonio Dias, Margo Seltzer, Michael D. Smith

February 1996 **ACM Transactions on Computer Systems (TOCS)**, Volume 14 Issue 1

**Publisher:** ACM Press

Full text available:  [pdf\(2.38 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This article presents a comparative study of the performance of three operating systems that run on the personal computer architecture derived from the IBM-PC. The operating systems, Windows for Workgroups, Windows NT, and NetBSD (a freely available variant of the UNIX operating system), cover a broad range of system functionality and user requirements, from a single-address-space model to full protection with preemptive multitasking. Our measurements are enabled by hardware counters in Intel ...

**Keywords:** Microsoft Windows, operating systems performance measurement, operating systems structure, personal computers

9 The impact of a zero-scan Internet checksumming mechanism

 Gregory G. Finn, Steve Hotz, Rod Van Meter  
October 1996 **ACM SIGCOMM Computer Communication Review**, Volume 26 Issue 5

Publisher: ACM Press

Full text available:  pdf(1.12 MB) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

This paper describes a "zero-scan" mechanism that reduces Internet checksumming overhead from a per-byte scan (or copy) cost, to a small and constant per-message cost. Unlike previous techniques, this mechanism requires no message buffering within the source. This will allow Internet transport protocols to achieve transfer latencies comparable to specialized protocols implemented directly on high-speed LAN (link-layer) services. In addition, this mechanism is transparent to systems outside of th ...

10 Architecture of the IBM system/370

 Richard P. Case, Andris Padegs  
January 1978 **Communications of the ACM**, Volume 21 Issue 1

Publisher: ACM Press

Full text available:  pdf(2.78 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper discusses the design considerations for the architectural extensions that distinguish System/370 from System/360. It comments on some experiences with the original objectives for System/360 and on the efforts to achieve them, and it describes the reasons and objectives for extending the architecture. It covers virtual storage, program control, data-manipulation instructions, timing facilities, multiprocessing, debugging and monitoring, error handling, and input/output operations. ...

**Keywords:** architecture, computer systems, error handling, instruction sets, virtual storage

11 Design choices in the SHRIMP system: an empirical study

 Matthias A. Blumrich, Richard D. Alpert, Yuqun Chen, Douglas W. Clark, Stefanos N. Damianakis, Cezary Dubnicki, Edward W. Felten, Liviu Iftode, Kai Li, Margaret Martonosi, Robert A. Shillner  
April 1998 **ACM SIGARCH Computer Architecture News , Proceedings of the 25th annual international symposium on Computer architecture ISCA '98**, Volume 26 Issue 3

Publisher: IEEE Computer Society, ACM Press

Full text available:  pdf(1.48 MB)  Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)  
[Publisher Site](#)

The SHRIMP cluster-computing system has progressed to a point of relative maturity; a variety of applications are running on a 16-node system. We have enough experience to understand what we did right and wrong in designing and building the system. In this paper we discuss some of the lessons we learned about computer architecture, and about the challenges involved in building a significant working system in an academic research environment. We evaluate significant design choices by modifying th ...

12 Performance prediction of a parallel simulator

Jason Liu, David Nicol, Brian J. Premore, Anna L. Poplawski

May 1999 **Proceedings of the thirteenth workshop on Parallel and distributed simulation**

Publisher: IEEE Computer Society

Full text available:  pdf(778.83 KB)  Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)  
[Publisher Site](#)

There are at least three major obstacles thwarting wide-spread adoption of parallel discrete-event simulation (a) lack of need, (b) lack of tools, (c) lack of predictability in behavior and performance. The plain truth is that most simulation studies can be adequately done on

ordinary serial computers. Parallel simulation tools are products of re-search efforts, and simply don't stand up to the demands of modern software engineering. The results of 20 years of research in parallel simulation rev ...

13 [Compilation: Vectorizing for a SIMD DSP architecture](#)

 Dorit Naishlos, Marina Biberstein, Shay Ben-David, Ayal Zaks

October 2003 **Proceedings of the 2003 international conference on Compilers, architecture and synthesis for embedded systems**

Publisher: ACM Press

Full text available:  pdf(301.45 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The Single Instruction Multiple Data (SIMD) model for finegrained parallelism was recently extended to support SIMD operations on disjoint vector elements. In this paper we demonstrate how SIMD (SIMD on disjoint data) supports effective vectorization of digital signal processing (DSP) benchmarks, by facilitating data reorganization and reuse. In particular we show that this model can be adopted by a compiler to achieve nearoptimal performance for important classes of kernels.

**Keywords:** SIMD, compiler controlled cache, data reuse, rotating register file, subword parallelism, vectorization, viterbi

14 [Learning not to share](#)

Jason Liu, David Nicol

May 2001 **Proceedings of the fifteenth workshop on Parallel and distributed simulation**

Publisher: IEEE Computer Society

Full text available:  pdf(779.87 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

 Publisher Site

*Strong reasons exist for executing a large-scale discrete-event simulation on a cluster of processor nodes (each of which may be a shared-memory multiprocessor or a uniprocessor). This is the architecture of the largest scale parallel machines, and so the largest simulation problems can only be solved this way. It is a common architecture even in less esoteric settings, and is suitable for memory-bound simulations. This paper describes our approach to porting the SSF simulation kernel to t ...*

15 [Mondrian memory protection](#)

 Emmett Witchel, Josh Cates, Krste Asanović

October 2002 **ACM SIGPLAN Notices , ACM SIGARCH Computer Architecture News , ACM SIGOPS Operating Systems Review , Proceedings of the 10th international conference on Architectural support for programming languages and operating systems ASPLOS-X**, Volume 37 , 30 , 36 Issue 10 , 5 , 5

Publisher: ACM Press

Full text available:  pdf(1.53 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

Mondrian memory protection (MMP) is a fine-grained protection scheme that allows multiple protection domains to flexibly share memory and export protected services. In contrast to earlier page-based systems, MMP allows arbitrary permissions control at the granularity of individual words. We use a compressed permissions table to reduce space overheads and employ two levels of permissions caching to reduce run-time overheads. The protection tables in our implementation add less than 9% overhead to ...

16 [Avoidance and suppression of compensation code in a trace scheduling compiler](#)

 Stefan M. Freudenberger, Thomas R. Gross, P. Geoffrey Lowney

July 1994 **ACM Transactions on Programming Languages and Systems (TOPLAS)**, Volume 16 Issue 4

Publisher: ACM Press

Full text available: [pdf\(3.58 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Trace scheduling is an optimization technique that selects a sequence of basic blocks as a trace and schedules the operations from the trace together. If an operation is moved across basic block boundaries, one or more compensation copies may be required in the off-trace code. This article discusses the generation of compensation code in a trace scheduling compiler and presents techniques for limiting the amount of compensation code: avoidance (restricting code motion so that no compensatio ...

**Keywords:** SPEC89, instruction-level parallelism, performance evaluation, trace scheduling

**17 Software support for outboard buffering and checksumming**

 Karl Kleinpaste, Peter Steenkiste, Brian Zill

October 1995 **ACM SIGCOMM Computer Communication Review , Proceedings of the conference on Applications, technologies, architectures, and protocols for computer communication SIGCOMM '95**, Volume 25 Issue 4

Publisher: ACM Press

Full text available: [pdf\(1.22 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Data copying and checksumming are the most expensive operations when doing high-bandwidth network IO over a high-speed network. Under some conditions, outboard buffering and checksumming can eliminate accesses to the data, thus making communication less expensive and faster. One of the scenarios in which outboard buffering pays off is the common case of applications accessing the network using the Berkeley sockets interface and the Internet protocol stack. In this paper we describe the changes t ...

**18 Reviewed articles: Measuring the evolution of transport protocols in the internet**

 Alberto Medina, Mark Allman, Sally Floyd

April 2005 **ACM SIGCOMM Computer Communication Review**, Volume 35 Issue 2

Publisher: ACM Press

Full text available: [pdf\(1.48 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

In this paper we explore the evolution of both the Internet's most heavily used transport protocol, TCP, and the current network environment with respect to how the network's evolution ultimately impacts end-to-end protocols. The traditional end-to-end assumptions about the Internet are increasingly challenged by the introduction of intermediary network elements (middleboxes) that intentionally or unintentionally prevent or alter the behavior of end-to-end communications. This paper provides mea ...

**Keywords:** Internet, TCP, evolution, middleboxes

**19 Towards a theory of cache-efficient algorithms**

 Sandeep Sen, Siddhartha Chatterjee, Neeraj Dumir

November 2002 **Journal of the ACM (JACM)**, Volume 49 Issue 6

Publisher: ACM Press

Full text available: [pdf\(273.41 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We present a model that enables us to analyze the running time of an algorithm on a computer with a memory hierarchy with limited associativity, in terms of various cache parameters. Our cache model, an extension of Aggarwal and Vitter's I/O model, enables us to establish useful relationships between the cache complexity and the I/O complexity of computations. As a corollary, we obtain cache-efficient algorithms in the single-level cache model for fundamental problems like sorting, FFT, and an i ...

**Keywords:** Hierarchical memory, I/O complexity, lower bound

**20 Novel approaches: High-speed I/O: the operating system as a signalling mechanism**

 Matthew Burnside, Angelos D. Keromytis

August 2003 **Proceedings of the ACM SIGCOMM workshop on Network-I/O convergence: experience, lessons, implications**

Publisher: ACM Press

Full text available:  [pdf\(127.65 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

The design of modern operating systems is based around the concept of memory as a cache for data that flows between applications, storage, and I/O devices. With the increasing disparity between I/O bandwidth and CPU performance, this architecture exposes the processor and memory subsystems as the bottlenecks to system performance. Furthermore, this design does not easily lend itself to exploitation of new capabilities in peripheral devices, such as programmable network cards or special-purpose h ...

**Keywords:** Architecture, Data Streaming, Operating Systems

Results 1 - 20 of 42

Result page: [1](#) [2](#) [3](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	9	(Herman near2 Dierks).in.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/02 10:20
L2	8	(Binh near2 Hua).in.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/02 10:20
L3	14	(Sivarama near2 Kodukula).in.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/02 10:21
L4	2	1 and 2 and 3	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/02 10:21
L5	22	1 or 2 or 3	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/02 10:21
L6	2626	address near2 align\$6	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/02 10:21
L7	2461	address near2 boundary	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/02 10:22
L8	2437	cross adj2 boundary	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/02 10:22
L9	195093	data adj transfer\$4	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/02 10:22
L10	145958	("end" or "near") near2 frame\$2	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/02 10:23

L11	7269	network adj adapter	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/02 10:23
L12	717	cache near3 offset	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/02 10:23
L13	4702	6 or 7	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/02 10:24
L14	2123	13 and 9	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/02 10:24
L16	3	15 and 10	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/02 10:24
L17	0	16 and 12	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/02 10:24